RISK ASSESSMENT

For

Central Web MediaWiki SaaS

Minor Application

CS-doc-6514

Responsibility for the Minor Application and its operation as described in this plan is accepted by:

_____ Date:_____

Service Owner - Peter J. Rzeminski II

_____ Date:_____

ESO Department Head – Michael Rosier

_____ Date:_____

Chief Information Security Officer – Irwin Gaines

_____ Date:_____

CCD Division Head – Jon Bakken

Contents

Information and Security Contacts:

Title	**Name**	**email**	**Telephone**	**Initials**
Service Owner (s)	Peter J. Rzeminski II	ptr@fnal.gov	630.840.5524	
System Managers – Apache httpd	Andrew Duranceau John Inkmann	adurance@fnal.gov inkmann@fnal.gov	630.840.6457 630.840.6508	
System Managers – NAS	Andrew Romero	romero@fnal.gov	630.840.4733	
System Managers – Linux OS	James O'Leary	joearly@fnal.gov	630.840.2230	
System Managers – Virtual Environment	Briant Lawson	blawson@fnal.gov	630.840.2944	
Management Contact	Jon Bakken (Division) Michael Rosier (Dept) Peter J. Rzeminski II (Group)	bakken@fnal.gov mrosier@fnal.gov ptr@fnal.gov	630.840.4790 630.840.8385 630.840.5524	

SYSTEM IDENTIFICATION

System Name/Title

Fermilab identifier CSP- GSS-#### has been assigned to the system discussed throughout this Risk Assessment and will be referred to as the Central Web MediaWiki SaaS.

System Type

This system is the Central Web MediaWiki SaaS Minor Application (MA) and is contained in the General Computing Enclave.

Responsible Organization

Fermi National Accelerator Laboratory
PO Box 500
Batavia, IL 60510

System Operational Status

It is in the Operational phase of its life-cycle.

General Description/Purpose

The Central Web MediaWiki SaaS is a centrally managed Software-as-a-Service providing MediaWiki website instances from a shared code base. Its purpose is to provide wiki-style websites for all employees, groups, offices, departments, experiments and any other approved entity connected to Fermilab without the need for the website owner to manage the software.

Introduction

The layout of the web infrastructure is designed to provide maximum stability and uptime for minimal cost. The service uses a pair of virtual servers running Red Hat Enterprise Linux (RHEL) behind the F5 Big Iron load balancer. Content is stored on the BlueArc and is mounted via NFSv4 to each web server.

The main components of the web service is the F5 Big Iron (F5), the web servers, and the BlueArc.

The flow of web traffic is broken up into a number of layers with each layer designed to provide as much redundancy as possible.

The initial layer is the F5 Big Iron (F5) where web traffic flows in the F5 and is directed to the second layer; the web servers.

The web servers are three components, the active server, the stand-by server, and the Site-Down Service servers. If the active server is unavailable, traffic is directed to the stand-by server. If both the active and the stand-by servers are unavailable, then traffic is directed to the Site-Down Service.

When traffic reaches the web servers, the Apache httpd process sets up an SSO session (optionally for read requests, necessary for write requests) and routes the request to the MediaWiki software running on PHP-FPM.

Security is maintained by adhering to the given baselines to each layer and by the conservative application of ACL's to the files and directories on the content file system. Access is solely through the web interface and is managed by utilizing existing authentication infrastructure; specifically the SERVICES AD realm. Access is granted/removed by the website owners through a web GUI tool like the one used for WordPress.

This is the Tier 6 category of the Central Web Service. It is similar to Tier 5 (WordPress), except there is no split between port 80 and port 443. Only port 443 is used. Users will NOT be allowed to SSH into the webserver.

Authentication

MediaWiki is only accessed through the web browser including website owners, content editors, and content viewers. It is authenticated via SSO/PingFederate using SERVICES Domain credentials.

System Description and Boundaries

The central web service infrastructure relies on the following third-party services

- Big Iron F5
    - Managed by the Networking Department, it is used to load-balance inbound traffic
- Virtual Server Infrastructure
    - Managed by the ESO Department / VMS Group, it is used to run the Apache httpd servers
- Red Hat Enterprise Linux (RHEL)
    - Managed by the ESO Department / USS Group, this is the operating system used for the service.
- Apache httpd Software
    - Managed by the ESO Department / WSA Group, this is the software used to present web content to the Internet.
- MediaWiki software
    - Managed by the ESO Department / WSA Group, this is the PHP software used to provide the wiki functionality.
- ElasticSearch software
    - Managed by the ESO Department / WSA Group, this is the software used to provide the back-end search engine for searching MediaWiki content.

- BlueArc File Server (NAS)
    - Managed by the ESO Department, this is the NFS File mount file system where user content and configuration files are stored.
- User Content
    - Managed by the customer & user, all content stored within the website is the responsibility of the site managers for the given website.

Virtual Server Design

System Name	Service Tier	Function	Virtual Server Location
Web600#	Tier VI	Standard Service offering for the customers needing a wiki website. See the Web Services SLA for the specifics.	FCC2

Webserver Cluster Design (Real World Example)

Cluster Name	Web Servers	Load Balancer Priority	Comments
VIP-WEBT6C01	web6001 131.225.70.49	150	Primary Content Server
	web6002 131.225.70.50	150	Secondary Content Server
	web-sorry01 131.225.70.234	50	Site Down Service – Server #1
	web-sorry02 131.225.70.14	1	Site Down Service – Server #2

Server Tier Offerings (non-exhaustive)

Offering	TIER VII
HTML	YES
JAVA	NO
PERL	NO
PHP	YES
PYTHON	NO
SAML/SSO Authentication to SERVICES via PingFederate	YES
MySQL/Postgres Client Libraries for Perl/PHP/Python etc.	YES
Direct SSH to Content	NO
Any MediaWiki extension	Upon request with justification and matching criteria

User Access Points

Method	Availability	Port(s)	Source	Access Requirements
HTTPS	World Wide	443	Apache httpd web server	None
SAML/SSO	World Wide	80/443	PingFederate	Optional secure authentication method through the browser
SSH **SCP** **SFTP**	Fermilab Network Only	22	Apache httpd web server	Available to Web Systems Administration Group only.
SMB, CIFS, Windows Share	None	139 & 445	BlueArc NAS	Not Available

Internet Traffic Flow

Central Web Service
Internet Traffic Flow

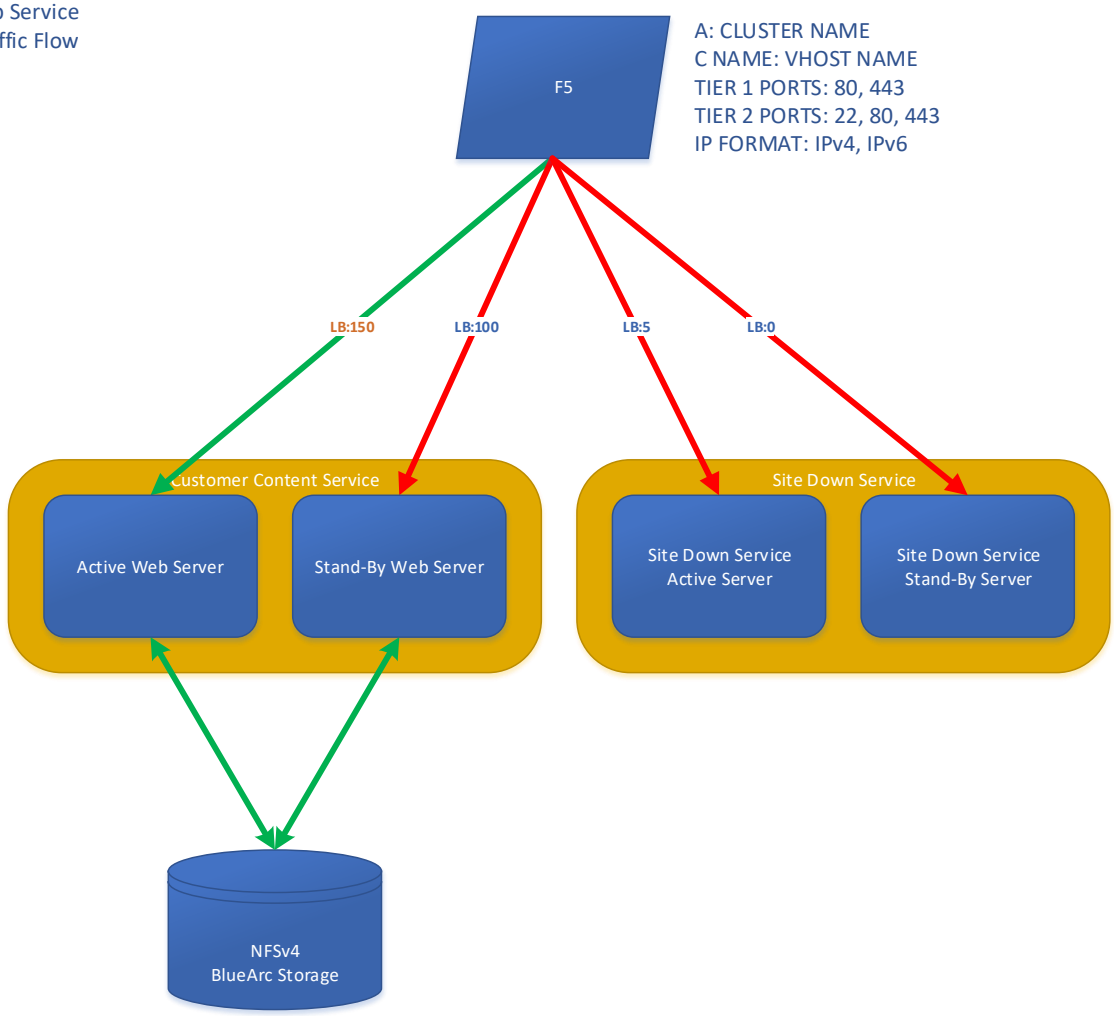

Internet Traffic Flow: User Content Access Methods

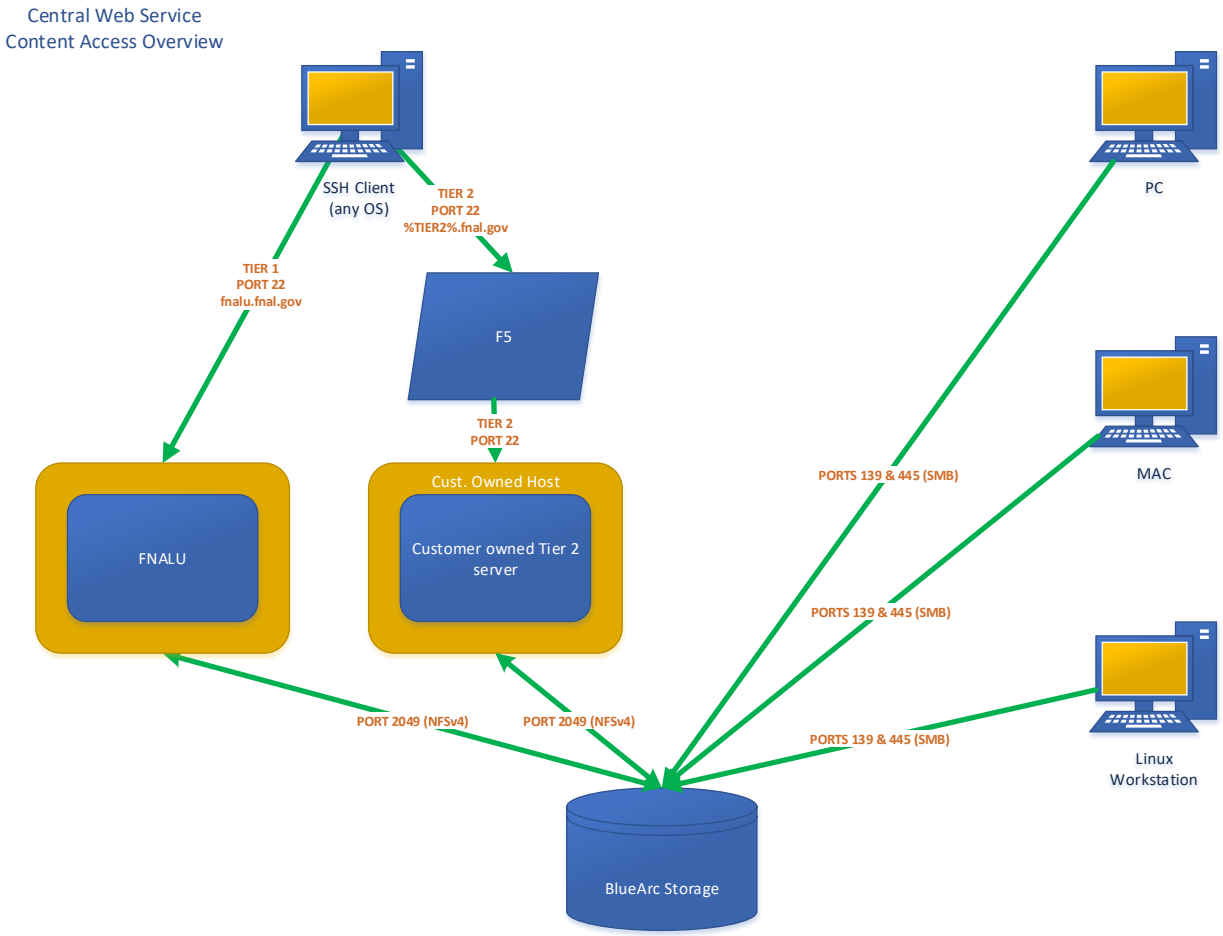

Server Startup: Configuration File Processing & Load Order

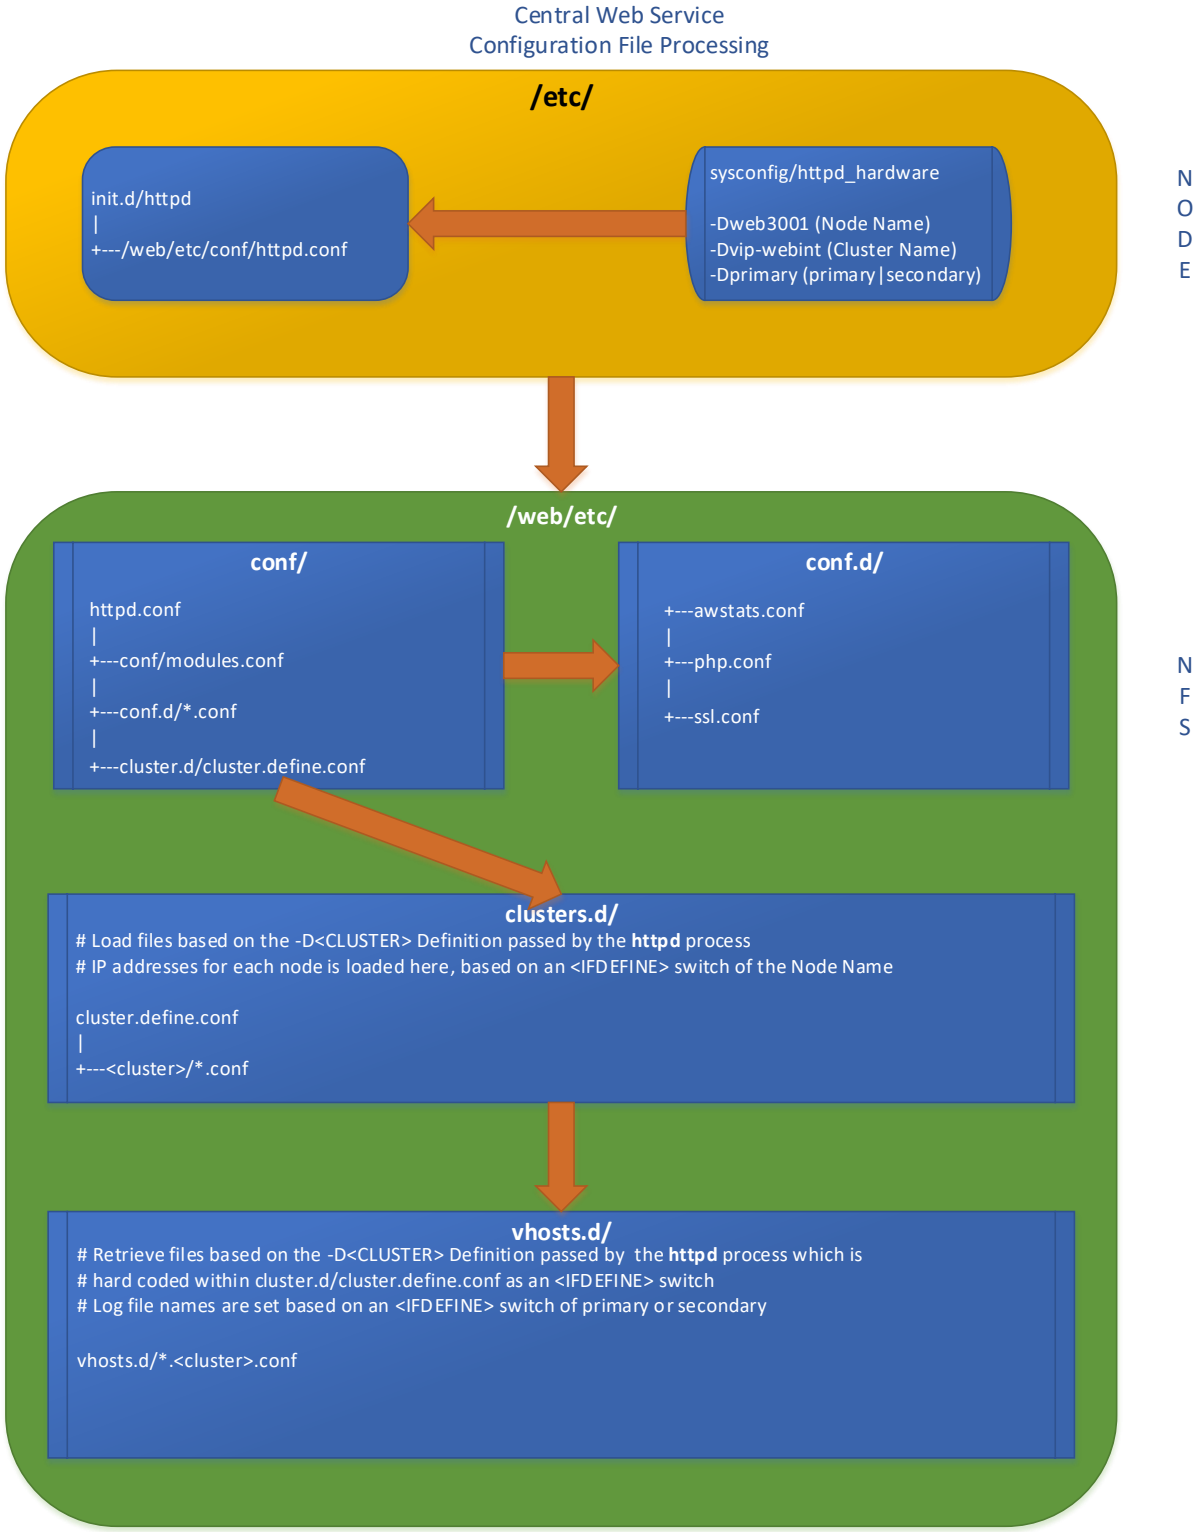

NFSv4 ACL Example

Central Web Service
NFSv4 ACL Overview

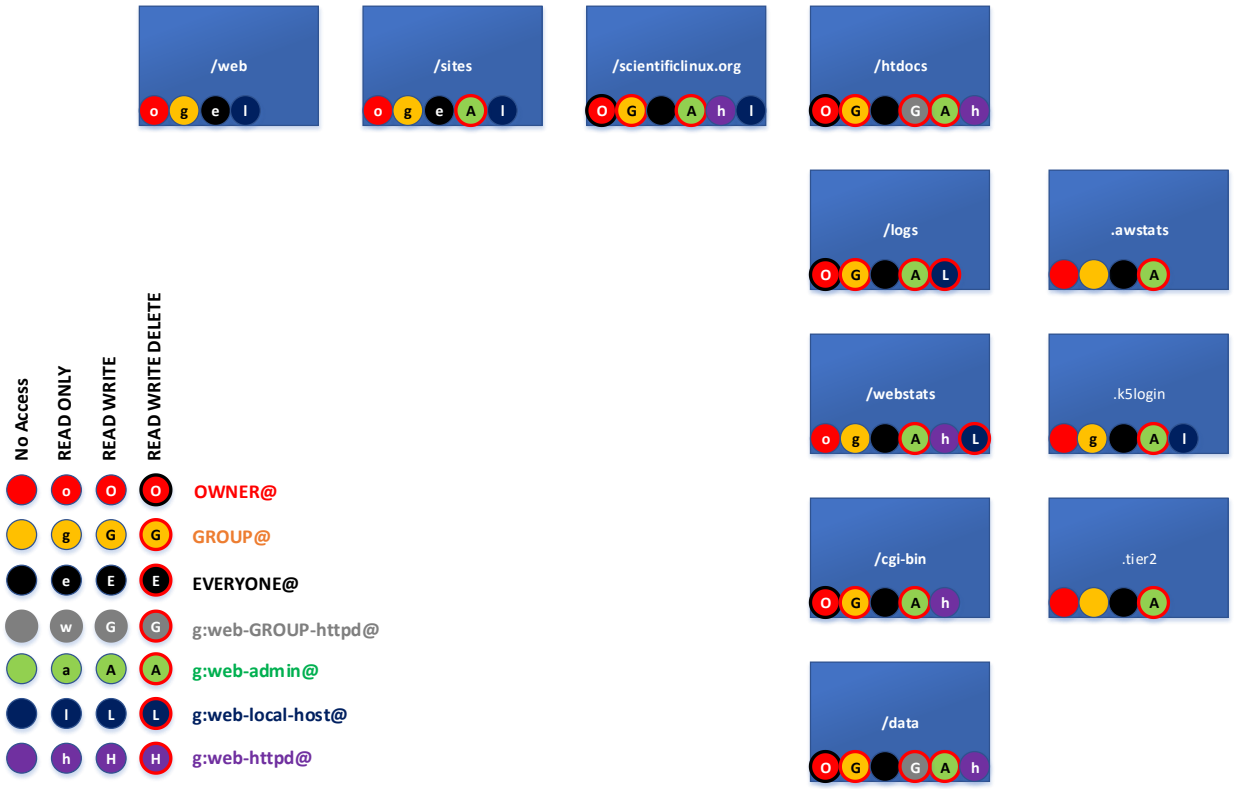

Subnet Allocation

All web servers will be on the 70 subnet. This subnet is reserved exclusively for use by Central Web Hosting.

VLAN ID	Network
VLAN 70 WWW-SUBNET	131.225.70.0/24 2620:6a:0:70::70:0/24
	gw: 131.225.70.1 gw: 2620:6a:0:70::70:1

Web Server Logs

Per standing Computer Security policy, all web server logs are being forwarded to clogger.fnal.gov.

Patching

Patching of the web service is defined in detail in the OLA between the WSA & LSS Groups. In brief, it will be handled in this manner:

- All patches that are of an urgent nature, as identified by the Vendor and/or Computer Security, will be applied immediately.
- All general patches supplied by the vendor will be applied on a monthly basis; Apache httpd, Perl, Python, and PHP are excluded from this general patching cycle and held back a few days to allow the Service Owners to test them on the Tier 3 dev & integration servers.
- Patches to the MediaWiki software will be done quarterly in accordance with vendor releases.
- When all patches have been cleared for use, the will be applied to the stand-by server first, the server checked to ensure to functions as expected, then applied to the active server.

NFSv4 ACL's

Content access is managed by the use of NFSv4 ACL's. Whereas a POSIX ACL is designed to be one dimensional in that the permissions are fixed to applied to only the owner, the group, and everybody else, NFSv4 ACL's are more three-dimensional. By that we mean that a single file or directory can have multiple owners, group, or hosts that can be granted or denied access, with each of them potentially being different types of access.

The ACL mapping is described like this:

```
                          A:fdnig:ENTITY@:rwaDdxtTnNcCoy
                          |||||||     |   :|||||||||||||
ACE Type (A)llow or (D)eny +:|||||     |   :|||||||||||||+ s(y)nchronize
            (f)ile inherit --+||||     |   :||||||||||||+- change (o)wner
      (d)directory inherit ---+|||     |   :|||||||||||+-- write a(C)l
    (n)o propagate inherit ----+||     |   :||||||||||+--- read a(c)l
            (i)nherit only -----+|     |   :|||||||||+---- write (N)amed attrib
       ENTITY is a (g)roup ------+     |   :||||||||+----- read (n)amed attrib
                                       |   :|||||||+------ write a(T)trib
This is the user that has  ------------+   :||||||+------- read a(t)trib
access to the file or                      :|||||+-------- e(x)ecute
directory.  If the (g)                     :||||+--------- (d)elete
flag above is listed                       :|||+---------- (D)elete child (Directory only)
then the entity is a                       :||+----------- (a)ppend / create-subdirectory
group and tied to a NAS                    :|+------------ (w)rite data / create-file
group and the users therein                :+------------- (r)read data  / list-directory
```

Note: ACL's are "default deny", which means if it is not explicitly granted access via an "A" type ACL, then the entity trying to access the file is denied access. While it is possible to create a "D" type ACL, it is frowned upon as different OS's may interpret the ACL differently. As a result, all ACL's created and used within the NAS are of the "A" type and configured to be least-permissive so that only those entities that need access are granted access.

Example ACL's

The ACL for the root directory of a web vhost looks like this:

```
A:fd:OWNER@:rwaDdxtncy
A:fd:GROUP@:rwaDdxtncy
A:fd:EVERYONE@:tcy
A:fdg:web-admin@fnal.gov:rwaDdxtTnNcCoy
A:fdng:web-httpd@fnal.gov:rxtncy
A:fdng:web-local-host@fnal.gov:rxtncy
```

- The owner of the directory has full rights.
- The group which owns the content, has the same rights.
- Anonymous users (everyone) has no rights to even see the content.

- web-admin, which is a group consisting of members of the ESO/WSA Group, have full rights to the content.
- The apache user, identified by the web-httpd group, has read-only access to the directory.
- The root user, identified by the web-local-host group, has read-only access to the directory.

The ACL for the logs directory under the directory above looks like this:

```
A:fd:OWNER@:rwaDxtTnNcy
A:fd:GROUP@:rwaDdxtTnNcy
A:fd:EVERYONE@:tcy
A:fdg:web-admin@fnal.gov:rwaDdxtTnNcCoy
A:fdg:web-local-host@fnal.gov:rwadxtTnNcy
```

- The owner of the directory has full rights.
- The group which owns the content, has the same rights.
- Anonymous users (everyone) has no rights to even see the content.
- web-admin, which is a group consisting of members of the ESO/WSA Group, have full rights to the content.
- The apache user, identified by the web-httpd group, is not listed in this ACL, meaning that it defaults to use the EVERYONE@ ACL, which blocks it access to the directory.
- The root user, identified by the g:web-local-host group, has read/write/delete access to the directory and the files within.

Information Sensitivity

The data sensitivity on ID is classified in the following table:

Relative Importance of Protection Needs			
	HIGH (Critical Concern)	**MEDIUM (Important Concern)**	**LOW (Minimum Concern)**
Confidentiality			**X**
Integrity			**X**
Availability			**X**

Risk Identification & Methodology

A threat is the potential for a particular threat-source to successfully exercise a particular vulnerability.

A vulnerability is a weakness that can be accidentally triggered or intentionally exploited.

A threat-source does not present a risk when there is no vulnerability that can be exercised.

Likelihood Determination, Impact Analysis, and Risk Level

The likelihood that each vulnerability will be exploited and the impact of a successful exploit is indicated by the pair of rankings associated with each vulnerability below. Following each vulnerability is the risk level obtained by using the following matrix:

		Impact		
		Low	**Medium**	**High**
Threat Likelihood	**Low**	Low	Low	Low
	Medium	Low	Medium	Medium
	High	Low	Medium	High

Threat Source Identification

There are no threat sources which have not been identified in the Risk Assessment for the General Computing Enclave.

Motivation and Threat Actions

There are no motivations and threat actions which have not been identified in the Risk Assessment for the General Computing Enclave.

Residual Risk Definition

Residual risks are divided into categories based on expected frequency of occurrence after full implementation of all security controls. We consider an occurrence rate to be:

LOW if it is expected to happen <10 times per year
VERY LOW if it is expected to happen <1 time per year
EXTREMELY LOW if it is expected to happen <1 time per five years

Identified Risks, Mitigations, and Residual Risks

General Risks

MediaWiki SaaS is an Enhanced Offering of the Central Web Hosting Service, built upon Apache httpd as the base service. As such, it is understood that most, if not all, of the Identified Risks, Mitigations, and Residual Risks covered by the Central Web Hosting Risk Assessment found in DocDB Doc.# 5346 will apply here. For the sake of brevity, this document will not repeat those Risks in this document. Instead it will document new Risks, or in the case where a CWH Risk has significantly changed due to the specific configuration differences of the MediaWiki SaaS offering, identify them here as something new with a reference to the original.

Risk: Hacking MediaWiki Software (General)

Threat & Vulnerability: Attackers will see that we are running a MediaWiki site and make every possible attempt to hack the software, take ownership of it, and use it for their own purposes. That could be anything from turning us into a spam relay, a node for a DDoS attack, scanning our internal processes, keylogging, to simply hosting their own content on our servers.

Mitigation: We have taken a great many precautions to protect the MediaWiki software from outside attack. By default, all MediaWiki websites are restricted by Apache to the Fermilab Intranet and only approved websites are unlocked for public access. Each MediaWiki instance is configured with unique database credentials. All database write functions are disabled unless the user has a valid SSO session to begin with, before individual authorizations are processed. The content directories are given Tier 1 permissions, meaning the webserver cannot write to publicly served directories.

Residual Risk: **Extremely Low** – We recognize that not all Risks can be eliminated when running software as a platform, but we believe we have taken every precaution possible to protect against external attackers.

Authentication Risks

Risk: Group Passwords

Since some users of MediaWiki websites (DUNE collaborators as of this writing) do not have SERVICES credentials, it will be possible to enable a built-in group password on a case basis. This creates a risk of unauthorized access to protected content as this password will be shared between multiple users.

Mitigation: The mitigation is not so much in what actions are being taken to protect the website, but the limited use of the password on the website itself. This password will only be used for read access and is impossible to use for write access due to the way the authorization is configured. As it is expected that this password will only be used on websites where lightweight protection is required, we do not see a great risk from its use.

Residual Risk: **LOW** – Using a group password means that the likelihood of unknown persons gaining access to it is much greater that with individually assigned credentials. However, the overall risk is insignificant because the content stored in those pages is unlikely to be anything that needs serious protection.

Service Specific Configuration Risks

Risk: MediaWiki Patching

Threat & Vulnerability: Patching of the MediaWiki software, which includes themes, plugins, and the core software, is managed by the WSA Group. As patches are released, it is possible that one of the software packages could break the plugin, theme, or core software that is already configured and running.

Mitigations: Patches will be applied to test websites first. These websites will be reviewed before and after the patch to determine if any detrimental behavior has occurred due to the new software. All patches will be documented through the Change Management process.

Residual Risk: **Extremely Low** – It is possible that a patch may be released that will not display any immediate server problems until specific conditions occur. The historical occurrence of this type of situation is near zero, making the likelihood extremely low.

Risk: User Access Points

Threat & Vulnerability: There is a threat from users directly accessing the content of their website on the file system and accidently deleting something and breaking the website., or directly modifying code, making it vulnerable to outside attack.

Mitigation: We have a single, fool-proof mitigation to users accessing content and breaking something; we do not let users directly access their content. The file system, across all protocols in all instances and occasions is off-limits to everybody except the Central Web Hosting group, and they will only access it when configured a website for use the first time.

Residual Risk: **Extremely Low** – In fact, essentially zero.

Risk: Third-Party Plugins & Themes

Threat & Vulnerability: MediaWiki sites can function without the use of Plugins and additional Themes, however, the Site Owners will demand their use to improve the functionality of their website. This is one area that we might carefully control to ensure we do not end up with the situation we had in the past with Plone, where the core software could not be updated due to plugin incompatibilities.

Mitigation: The mitigations will be unpopular, but are necessary to ensure the long-term health and security of the MediaWiki SaaS offering as a whole. They are the same mitigations used for WordPress.

- No Plugin or Theme will be installed to a Production MediaWiki SaaS website without being approved by the Central Web Hosting Service Owner. The approval process will be documented and published publically to ensure it is conducted in neutral and non-arbritary a manner as possible.
- To be approved for installation, it must pass a set of criteria that will be documented and published in DocDB. This document will be reviewed on a regular basis and submitted for peer-review before being annually updated.
- While the full list of criteria is still being documented, the following items are already known:
    - The plugin must not modify security features or permissions.
    - The software must work with the InnoDB database engine being used MariaDB.
    - The software must be packaged as a single file and not require installation steps beyond running MediaWiki's standard maintenance script.
- In addition, the following is a non-exhaustive list of how Plugins and Themes will be treated during their lifespan on the MediaWiki SaaS offering.
    - If a Plugin or Theme becomes a security risk due to problems in its code or how it interacts with the Fermilab infrastructure, it will be uninstalled across all websites. Site Owners will be notified of its removal, in-advance whenever possible, posthumously if the need to remove it was great.
    - The ability to upgrade the OS and the MediaWiki software is paramount over any Plugin & Theme. If MediaWiki releases a new version of their core software and some Plugins & Themes become unusable on the new version of the software, said Plugins & Themes will be disabled in order to upgrade the core software. If the authors of the Plugins & Themes do not upgrade their own software to make it compatible with the new core MediaWiki software, then they will be uninstalled after a reasonable period of time.

Residual Risk: **Low** – Most plugins are compatible across core software versions, and if they are broken, have historically been updated shortly after the core software update. However, the possibility of plugins being disabled by breaking changes in the core software is more likely than with WordPress due to the MediaWiki plugin developer community being less organized. Therefore, this is stated as a Low risk instead of Very Low.

Risk: Information disclosure through malicious search requests

Threat & Vulnerability: MediaWiki's built-in search engine is not compatible with MariaDB & Galera. As an alternative, we have configured MediaWiki to submit article text to be indexed by an ElasticSearch server and to query that server for search requests. Each MediaWiki site uses its own index table, but there is no granular access control provided by ElasticSearch to protect against crossing the boundary between indices. In theory, it is possible for a user to view information stored in an index belonging to another wiki site by submitting a maliciously crafted search request.

Mitigation: If it is possible to access information from another search index, it would be limited to indices stored in the same ElasticSearch instance. We have configured multiple ElasticSearch instances to limit this risk to similarly controlled information:

- Public instance
  Stores indices for websites that can be viewed by the public without authentication

- Internal instance
  Stores indices for websites that can only be viewed by users with SERVICES credentials

- Private instance
  Stores indices for websites that can only be viewed by users in a specific group

An authenticated Apache proxy is placed in front of each ElasticSearch instance with its own set of credentials. These credentials are shared only between websites with the same level of access control. Firewall rules are in place so that only the MediaWiki webservers can access this proxy, and only Central Web administrators can access ElasticSearch directly.

Residual Risk: Extremely Low – We are not aware of any attacks of this nature; it is only a theoretical possibility. MediaWiki websites are only to be used for storing information of Low sensitivity. The separation of instances according to website access level further mitigates the level of risk.